

What is claimed:

1. In a communication network comprising:

at least one local name server, each said local name server being capable of answering name-to-address resolution queries by using temporarily stored information or by further querying other name servers,

at least one application server, each said application server having at least one application server address and being capable of receiving requests for at least one service and of performing said service, wherein each said service has a distinct service name,

a plurality of clients, each said client being associated with at least one said local name server, and being capable to query any said associated local name server for the address of an application server providing a service with a specified service name, to receive from said associated local name server an answer specifying an address of one said application server, and to send a request for that service to that application server,

and at least one authoritative name server, each said authoritative name server being capable of answering name-to-address resolution queries from said local name servers about any of said services, the content of every answer having a validity period,

a method for discovering associations between clients and local name servers, comprising the steps of:

(A) producing a query record concerning a query received by an authoritative name server and the answer to that query, said query record comprising at least some of the following information items:

(q1) an identifier of the application server that is the answer to the query,

(q2) a name of a service provided by the application server,

(q3) a timestamp expressing the moment in time at which the answer is issued,

(q4) a validity period defined for the identifier of said application server address, and

(q5) an address of the local name server from which the query is received,

wherein said item (q1) is mandatory if for all said application servers the union of their respective sets of application server addresses comprises more than one member; said item (q2) is optional; said item (q3) is mandatory; said item (q4) is mandatory if said validity period may differ from one said answer to another, otherwise it is a predetermined constant value; and said item (q5) is mandatory;

(B) producing a request record concerning a service request received by an application server, said request record comprising at least some of the following information items:

(r1) an identifier of the application server at which the request is received,

(r2) a name of the service provided by the application server,

(r3) a timestamp expressing the moment of time at which the request is received, and

(r4) an address of the client having issued the request;

wherein said item (r1) is mandatory if for all said application servers the union of their respective sets of application server addresses comprises more than one member; said item (r2) is optional, and said items (r3) and (r4) are mandatory; and

(C) finding matching pairs of one said query record and one said request record, and associating the address of the local name server from the query record in a matching pair to the address of the client from the request record in the same matching pair, a matching pair being defined as such first pair of one query record and one request record which satisfies the following conditions:

an address match condition that requires that the identifier of the application server be the same in the query record and in the request record of said first pair;

a timestamp match condition that requires the timestamp in the request record of said first pair to express a moment of time that is within said validity period starting at the moment of time expressed by the timestamp in the query record of said first pair,

a service name match condition that requires that the service name from the query record match the service name from the request record whenever both said records contain said service name item; otherwise, if at least one among said query record and said request record does not contain said service name item, then said service name match condition is considered to be satisfied; and

a uniqueness condition consisting in that no second pair of one query record and one request record be found to satisfy every said condition (1), (2) and (3) and to have the request record the same as the request record of the first pair, but the query record different from the query record of the first pair.

2. The method as recited in claim 1, further comprising the following dispositions ensuring that said uniqueness condition is satisfied for every said matching pair of records:

(D) selecting a set of monitored services as a non-empty subset of said services;

(E) for every said monitored service, selecting a set of monitored servers as a non-empty subset of all application servers capable of performing said monitored service, each said monitored server having at least two distinct application server addresses;

(F) for every said monitored server, dividing its application server addresses into two disjoint and non-empty subsets, a first set of monitoring addresses for each of which said request records are to be produced, and a second set of standard addresses, for each of which said request records are not to be produced;

(G) considering any said monitoring address as engaged in a given service since the moment when some authoritative name server issues that monitoring address in an answer to a query for said service, and until the moment when the validity period of said answer expires, and considering said monitoring address as free from said service at any time when it is not engaged in said service, and further considering said monitoring address absolutely free if it is free from any service; and

(H) instructing an authoritative name server, when it receives a name-to-address resolution query about one specific monitored service to select an address of one specific monitored server in the following way:

(1) if said option of registering the queried service name (q2) in said query records is chosen at said authoritative name server, and said option of registering the requested service name (r2) in said request records is chosen at said monitored server, then to take said address from the list of monitoring addresses of said monitored server that are free from said service, whenever said list is not empty;

(2) otherwise, to take said address from the list of absolutely free monitoring addresses of said monitored server, whenever said list is not empty;

(3) otherwise, when all said monitoring addresses are engaged, to use a standard address of said monitored server.

3. The method as recited in claim 2, further comprising the following step:

(I) for at least one local name server, and for at least some of queries received by an authoritative name server from that local name server and concerning a monitored service, instructing said authoritative name server to answer said queries in alternate fashion, by answering every said query, whenever possible, with a monitoring address of any one among those monitored servers that are capable of performing said monitored service and whose monitoring addresses were not used in the most recent answers of said authoritative name server to the same local name server.

4. The method as recited in claim 1, further comprising the steps of:

when discovering an association between a client and a local name server, producing an association record comprising the following information items:

(a1) an address of a client, and

(a2) an address of the local name server associated to said client;

given a predetermined set of communication parameters, producing a measurement record for a service request received by an application server, said measurement record comprising the following information items:

(m1) an application server address at which the request is received,

(m3) an address of the client having issued the request, and

(m3) a list of client-to-server values of at least some of communication parameters of said set of communication parameters, measured for communication of said client with said application server;

for a given application server and a given local name server, collecting all said measurement records related to said application server and to said local name server, using said association records;

for a given application server and a given local name server, calculating aggregate values for at least some of said communication parameters, by applying appropriate mathematical methods to client-to-server values of said parameters as taken from said measurement records, said aggregate values characterizing said local name server with regard to said application server.

5. The method as recited in claim 3, further comprising the steps of:

(J) when discovering an association between a client and a local name server, producing an association record comprising the following information items:

(a1) address of a client, and

(a2) address of the local name server associated to said client;

(K) given a predetermined set of communication parameters, producing a measurement record for a service request received by an application server, said measurement record comprising the following information items:

(m1) an application server address at which the request is received,
(m3) an address of the client having issued the request, and
(m3) a list of client-to-server values of at least some of communication parameters of said set of communication parameters, measured for communication of said client with said application server;

(L) for a given application server and a given local name server, collecting all said measurement records related to said application server and to said local name server, using said association records;

(M) for a given application server and a given local name server, calculating aggregate values for at least some of said communication parameters, by applying appropriate mathematical methods to client-to-server values of said parameters as taken from said measurement records, said aggregate values characterizing said local name server with regard to said application server.

6. The method as recited in claim 5, further comprising the following steps:

(N) for at least one local name server, and for at least one monitored service, repeating said step (I) a plurality of times to collect aggregate values of at least some of said communication parameters, characterizing said local name server with regard to a plurality of monitored servers;

(O) after said collecting said aggregate values, pausing performance of said step (I), and assigning to said local name server that monitored server considered as the best within said plurality of monitored servers according to some predefined criteria applied to aggregate values;

(P) after performing said assignment, answering subsequent queries from said local name server concerning said monitored service with addresses of said best monitored server.

7. The method as recited in claim 6, further comprising the following steps:

(Q) when performing said step (P), continuing to calculate said aggregate values characterizing said local name server with regard to said monitored server, by using more recent values of said communication parameters;

(R) whenever current aggregate values of communication parameters, characterizing said local name server with regard to said currently assigned monitored server, leave some predefined tolerance range, pausing performance of step (P), repeating performance of said steps (N) and (O), and restarting performance of step (P) using the most recently assigned best monitored server.

8. The method as recited in claim 6, further comprising the following step:

(S) after performing said step (P) for a predefined period of time, pausing its performance, repeating performance of said steps (N) and (O), which may result in re-assignment to said local name server of yet another better monitored server, and then restarting performance of step (P) using the most recently assigned best monitored server.

9. A system for discovering associations between clients and local name servers, comprising:

at least one name server monitor, every said name server monitor being associated with a particular authoritative name server and producing query records with information from queries to that authoritative name server and from answers to said queries;

at least one application server monitor, every said application server monitor being associated with a particular application server and producing request records with information from service requests to that application server;

and at least one discovery and monitoring manager, said discovery and monitoring manager collecting information gathered by at least some of said name

server monitors and by at least some of said application server monitors, and discovering associations between clients and local name servers.

10. The system of claim 9 wherein:

at least one said application server is accessible on two disjoint and non-empty sets of addresses, a first set of monitoring addresses for each of which said request records are to be produced by the associated application server monitor, and a second set of standard addresses, for which said request records are not to be produced;

at least one said name server monitor is able to force its associated authoritative name server to answer queries from local name servers with specified addresses of application servers;

and at least one discovery and monitoring manager also maintains a table of monitoring addresses currently engaged in said discovery process, and is able to instruct said name server monitors as to the currently free monitoring addresses to be supplied in answers to queries from local name servers undergoing said discovery process.

11. The system of claim 10 wherein:

at least one discovery and monitoring manager is able to instruct said name server monitors to force their associated authoritative name servers to answer successive queries from local name servers undergoing said discovery process with monitoring addresses of different application servers in alternate fashion.

12. The system of claim 9 wherein:

at least one application server monitor also measures values of communication parameters for communications between its associated application server and requesting clients;

and

at least one discovery and monitoring manager also calculates aggregate values of communication parameters characterizing local name servers with regard to application servers.

13. The system of claim 11 wherein:

at least one application server monitor also measures values of communication parameters for communications between its associated application server and requesting clients;

and

at least one discovery and monitoring manager also calculates aggregate values of communication parameters characterizing local name servers with regard to application servers, selects for a local name server an application server that is considered the best application server according to some predefined criteria applied to the aggregate values, assigns to said local name server the application server selected for it, and, after performing said assignment, instructs name server monitors to force their associated authoritative name servers to answer further queries from said local name server with a standard address of its assigned application server.

14. The system of claim 13, wherein:

at least one application server monitor continues measuring values of communication parameters for communications between its associated application server and requesting clients associated with a given local name server, even after that application server has been assigned to that local name server; and

at least one discovery and monitoring manager continues calculating aggregate values of said communication parameters, checks them against predefined tolerance range, repeats said procedure of selection and assignment of the best application server to a local name server whenever said aggregate values leave said tolerance range for that local name server, and, after performing new assignment, instructs name

server monitors to force authoritative name servers to answer further queries from that local name server with a standard address of its newly assigned application server.

15. The system of claim 13 wherein:

at least one discovery and monitoring manager also maintains a re-assignment time schedule for at least one local name server, performs a recurring procedure of selection and assignment of the best application server to that local name server according to said time schedule, and, after performing new assignment, instructs name server monitors to force authoritative name servers to answer further queries from that local name server with a standard address of its newly assigned application server.

16. In a communication network comprising:

at least one local name server, each said local name server being capable of answering name-to-address resolution queries by using temporarily stored information or by further querying other name servers,

at least one application server, each said application server having at least one application server address and being capable of receiving requests for at least one service and of performing said service, wherein each said service has a distinct service name,

a plurality of clients, each said client being associated with at least one said local name server and being capable to query an associated local name server for the address of an application server providing a service with a specified service name, to receive from said associated local name server an answer specifying an address of one said application server, and to send a request for that service to that application server,

and at least one authoritative name server, each said authoritative name server being capable of answering name-to-address resolution queries from said local name servers about any of said services, the content of every answer having a validity period,

a method for discovering associations between clients and local name servers, comprising the steps of:

producing a query record concerning a query received by an authoritative name server and the answer to that query, said query record comprising the following information items:

an identifier of the application server that is the answer to the query, a timestamp expressing the moment in time at which the answer is issued, and

an address of the local name server from which the query is received, producing a request record concerning a service request received by an application server, said request record comprising the following information items:

an identifier of the application server at which the request is received, a timestamp expressing the moment of time at which the request is received, and

the address of the client having issued the request; and

finding matching pairs of one said query record and one said request record, and associating the address of the local name server from the query record in a matching pair to the address of the client from the request record in the same matching pair, a matching pair being defined as a first pair of one query record and one request record which satisfies:

an address match condition that requires that the identifier of the application server be the same in the query record and in the request record; and

a timestamp match condition that requires the timestamp of the request record be within a predetermined time after the timestamp in the query record.

17. The method of claim 16 wherein the matching pair further satisfy:

a uniqueness condition, said uniqueness condition consisting in that no second pair of one query record and one request record be found to satisfy every condition in said set of match conditions and to have the request record the same as the request

record of the first pair, but the query record different from the query record of the first pair.

18. The method of claim 16 wherein the matching pair further satisfy:

a service name match condition which requires that a service name from the query record match a service name from the request record whenever both said records contain said service name item; otherwise, if at least one among said query record and said request record does not contain said service name item, then said service name match condition is considered to be satisfied.

19. The method of claim 16 further comprising the steps of:

producing a measurement record for a service request received by an application server, said measurement record comprising the following information items:

an application server address at which the request is received,
an address of the client having issued the request,

a list of client-to-server values for at least some communication parameters measured for communication of said client with said application server;

for a given application server and a given local name server, collecting all said measurement records related to said application server and to said local name server, using said associations between clients and local name servers; and

for a given application server and a given local name server, calculating aggregate values for at least some of said communication parameters, by applying mathematical methods to client-to-server values of said parameters as taken from said measurement records, said aggregate values characterizing said local name server with regard to said application server.

20. The method of claim 19, further comprising the step of:

assigning to said local name server that application server that is considered as the best application server according to some predefined criteria applied to the aggregate values.